

M5 simulator system
TDT4260 Computer Architecture
User documentation

Last modified: January 20, 2014

Contents

1	Introduction	2
1.1	Overview	2
1.2	Chapter outlines	2
2	Installing and running M5	4
2.1	Download	4
2.2	Installation	4
2.2.1	Linux	4
2.2.2	VirtualBox disk image	5
2.3	Build	5
2.4	Run	5
2.4.1	CPU2000 benchmark tests	6
2.4.2	Running M5 with custom test programs	7
2.5	Submitting the prefetcher for benchmarking	8
3	The prefetcher interface	10
3.1	Memory model	10
3.2	Interface specification	10
3.3	Using the interface	12
3.3.1	Example prefetcher	14
4	Statistics	15
5	Debugging the prefetcher	17
5.1	m5.debug and trace flags	17
5.2	GDB	17
5.3	Valgrind	18

Chapter 1

Introduction

You are now going to write your own hardware prefetcher, using a modified version of M5, an open-source hardware simulator system. This modified version presents a simplified interface to M5’s cache, allowing you to concentrate on a specific part of the memory hierarchy: a prefetcher for the second level (L2) cache.

1.1 Overview

This documentation covers the following:

- Installing and running the simulator
- Machine model and memory hierarchy
- Prefetcher interface specification
- Using the interface
- Testing and debugging the prefetcher on your local machine
- Submitting the prefetcher for benchmarking
- Statistics

1.2 Chapter outlines

The **first chapter** gives a short introduction, and contains an outline of the documentation.

The **second chapter** starts with the basics: how to install the M5 simulator. There are two possible ways to install and use it. The first is as a stand-alone VirtualBox disk-image, which requires the installation of VirtualBox. This is the best option for those who use Windows as their operating system of choice. For Linux enthusiasts, there is also the option of downloading a tarball, and installing a few required software packages.

The chapter then continues to walk you through the necessary steps to get M5 up and running: building from source, running with command-line options that enables prefetching, running local benchmarks, compiling and running custom test-programs, and finally, how to submit your prefetcher for testing on a computing cluster.

The **third chapter** gives an overview over the simulated system, and describes its memory model. There is also a detailed specification of the prefetcher interface, and tips on how to use it when writing your own prefetcher. It includes a very simple example prefetcher with extensive comments.

The **fourth chapter** contains definitions of the statistics used to quantitatively measure prefetchers.

The **fifth chapter** gives details on how to debug prefetchers using advanced tools such as GDB and Valgrind, and how to use trace-flags to get detailed debug printouts.

Chapter 2

Installing and running M5

2.1 Download

Download the modified M5 simulator from the **PfJudge**^β website.

2.2 Installation

2.2.1 Linux

Software requirements (specific Debian/Ubuntu packages mentioned in parentheses):

- 3.4.6 <= g++ <= 4.4
- Python and libpython >= 2.4 (`python` and `python-dev`)
- Scons > 0.98.1 (`scons`)
- SWIG >= 1.3.31 (`swig`)
- zlib (`zlib1g-dev`)
- m4 (`m4`)

To install all required packages in one go, issue instructions to `apt-get`:

```
sudo apt-get install g++ python-dev scons swig zlib1g-dev m4
```

The simulator framework comes packaged as a gzipped tarball. Start the adventure by unpacking with `tar xvzf framework.tar.gz`. This will create a directory named `framework`.

2.2.2 VirtualBox disk image

If you do not have convenient access to a Linux machine, you can download a virtual machine with M5 preconfigured. You can run the virtual machine with VirtualBox, which can be downloaded from <http://www.virtualbox.org>.

The virtual machine is available as a zip archive from the **PfJudge**^β website. After unpacking the archive, you can import the virtual machine into VirtualBox by selecting “Import Appliance” in the file menu and opening “Prefetcher framework.ovf”. The root password of the virtual machine is “user” (no quotation marks). It will be required if you want to run *sudo* commands.

2.3 Build

M5 uses the `scons` build system: `scons -j2 ./build/ALPHA_SE/m5.opt` builds the optimized version of the M5 binaries.

`-j2` specifies that the build process should built two targets in parallel. This is a useful option to cut down on compile time if your machine has several processors or cores.

The included build script `compile.sh` encapsulates the necessary build commands and options.

Since the `g++` version is rather old, you can also compile the simulator on kongull. Use `module load scons && module load swig` to load the necessary programs (check with `module list` whether they were loaded) and then compile by invoking the build script.

2.4 Run

Before running M5, it is necessary to specify the architecture and parameters for the simulated system. This is a nontrivial task in itself. Fortunately there is an easy way: use the included example python script for running M5 in syscall emulation mode, `m5/config/example/se.py`. When using a prefetcher with M5, this script needs some extra options, described in [Table 2.1](#).

For an overview of all possible options to `se.py`, do

```
./build/ALPHA_SE/m5.opt configs/example/se.py --help
```

When combining all these options, the command line will look something like this:

Option	Description
--detailed	Detailed timing simulation
--caches	Use caches
--l2cache	Use level two cache
--l2size=1MB	Level two cache size
--prefetcher=policy=proxy	Use the C-style prefetcher interface
--prefetcher=on_access=True	Have the cache notify the prefetcher on <i>all</i> accesses, both hits and misses
--cmd	The program (an Alpha binary) to run

Table 2.1: Basic `se.py` command line options.

```
./build/ALPHA_SE/m5.opt configs/example/se.py --detailed
--caches --l2cache --l2size=1MB --prefetcher=policy=proxy
--prefetcher=on_access=True
```

This command will run `se.py` with a default program, which prints out “Hello, world!” and exits. To run something more complicated, use the `--cmd` option to specify another program. See [subsection 2.4.2](#) about cross-compiling binaries for the Alpha architecture. Another possibility is to run a benchmark program, as described in the next section.

2.4.1 CPU2000 benchmark tests

The `test_prefetcher.py` script can be used to evaluate the performance of your prefetcher against the SPEC CPU2000 benchmarks. It runs a selected suite of CPU2000 tests with your prefetcher, and compares the results to some reference prefetchers.

The per-test statistics that M5 generates are written to `output/<testname-prefetcher>/stats.txt`. The statistics most relevant for hardware prefetching are then filtered and aggregated to a `stats.txt` file in the framework base directory.

See [chapter 4](#) for an explanation of the reported statistics.

Since programs often do some initialization and setup on startup, a sample from the start of a program run is unlikely to be representative for the whole program. It is therefore desirable to begin the performance tests after the program has been running for some time. To save simulation time, M5 can resume a program state from a previously stored checkpoint. The prefetcher framework comes with checkpoints for the CPU2000 benchmarks taken after 10^9 instructions.

It is often useful to run a specific test to reproduce a bug. To run the CPU2000 tests outside of `test_prefetcher.py`, you will need to set the `M5_CPU2000` environment variable. If this is set incorrectly, M5 will give the error message “Unable to find workload”. To export this as a shell variable, do

```
export M5_CPU2000=lib/cpu2000
```

Near the top of `test_prefetcher.py` there is a commented-out call to `dry_run()`. If this is uncommented, `test_prefetcher.py` will print the command line it would use to run each test. This will typically look like this:

```
m5/build/ALPHA_SE/m5.opt --remote-gdb-port=0 -re
--outdir=output/ammp-user m5/configs/example/se.py
--checkpoint-dir=lib/cp --checkpoint-restore=1000000000
--at-instruction --caches --l2cache --standard-switch
--warmup-insts=10000000 --max-inst=10000000 --l2size=1MB
--bench=ammp --prefetcher=on_access=true:policy=proxy
```

This uses some additional command line options, these are explained in [Table 2.2](#).

Option	Description
--bench=ammp	Run one of the SPEC CPU2000 benchmarks.
--checkpoint-dir=lib/cp	The directory where program checkpoints are stored.
--at-instruction	Restore at an instruction count.
--checkpoint-restore=n	The instruction count to restore at.
--standard-switch	Warm up caches with a simple CPU model, then switch to an advanced model to gather statistics.
--warmup-insts=n	Number of instructions to run warmup for.
--max-inst=n	Exit after running this number of instructions.

Table 2.2: Advanced `se.py` command line options.

2.4.2 Running M5 with custom test programs

If you wish to run your self-written test programs with M5, it is necessary to cross-compile them for the Alpha architecture. The easiest way to achieve this is to download the precompiled compiler-binaries provided by [crosstool](#) from the [M5 website](#). Install the one that fits your host machine best (32 or 64 bit version). When cross-compiling your test program, you must use the `-static` option to enforce static linkage.

To run the cross-compiled Alpha binary with M5, pass it to the script with the `--cmd` option. Example:

```
./build/ALPHA_SE/m5.opt configs/example/se.py --detailed  
--caches --l2cache --l2size=512kB --prefetcher=policy=proxy  
--prefetcher=on_access=True --cmd /path/to/testprogram
```

2.5 Submitting the prefetcher for benchmarking

First of all, you need a user account on the [PfJudge^β](#) web pages. The teaching assistant in TDT4260 Computer Architecture will create one for you. You must also be assigned to a group to submit prefetcher code or view earlier submissions.

Sign in with your username and password, then click “Submit prefetcher” in the menu. Select your prefetcher file, and optionally give the submission a name. This is the name that will be shown in the highscore list, so choose with care. If no name is given, it defaults to the name of the uploaded file. If you check “Email on complete”, you will receive an email when the results are ready. This could take some time, depending on the cluster’s current workload.

When you click “Submit”, a job will be sent to the [Kongull](#) cluster, which then compiles your prefetcher and runs it with a subset of the CPU2000 tests. You are then shown the “View submissions” page, with a list of all your submissions, the most recent at the top.

When the prefetcher is uploaded, the status is “Uploaded”. As soon as it is sent to the cluster, it changes to “Compiling”. If it compiles successfully, the status will be “Running”. If your prefetcher does not compile, status will be “Compile error”. Check “Compilation output” found under the detailed view.

When the results are ready, status will be “Completed”, and a score will be given. The highest scoring prefetcher for each group is listed on the highscore list, found under “Top prefetchers” in the menu. Click on the prefetcher name to go a more detailed view, with per-test output and statistics.

If the prefetcher crashes on some or all tests, status will be “Runtime error”. To locate the failed tests, check the detailed view. You can take a look at the output from the failed tests by clicking on the “output” link found after each test statistic.

To allow easier exploration of different prefetcher configurations, it is possible to submit several prefetchers at once, bundled into a zipped file. Each `.cc` file in the archive is submitted independently for testing on the cluster. The submission is named after the compressed source file, possibly prefixed with the name specified in the submission form.

There is a limit of 50 prefetchers per archive.

Chapter 3

The prefetcher interface

3.1 Memory model

The simulated architecture is loosely based on the DEC Alpha Tsunami system, specifically the [Alpha 21264](#) microprocessor. This is a superscalar, out-of-order (OoO) CPU which can reorder a large number of instructions, and do speculative execution.

The L1 prefetcher is split in a 32kB instruction cache, and a 64kB data cache. Each cache block is 64B. The L2 cache size is 1MB, also with a cache block size of 64B. The L2 prefetcher is notified on every access to the L2 cache, both hits and misses. There is no prefetching for the L1 cache.

The memory bus runs at 400MHz, is 64 bits wide, and has a latency of 30ns.

3.2 Interface specification

The interface the prefetcher will use is defined in a header file located at `prefetcher/interface.hh`. To use the prefetcher interface, you should include `interface.hh` by putting the line `#include "interface.hh"` at the top of your source file.

<code>#define</code>	Value	Description
<code>BLOCK_SIZE</code>	64	Size of cache blocks (cache lines) in bytes
<code>MAX_QUEUE_SIZE</code>	100	Maximum number of pending prefetch requests
<code>MAX_PHYS_MEM_SIZE</code>	$2^{28} - 1$	The largest possible physical memory address

Table 3.1: Interface `#defines`.

NOTE: All interface functions that take an address as a parameter block-align the address before issuing requests to the cache.

Function	Description
<code>void prefetch_init(void)</code>	Called before any memory access to let the prefetcher initialize its data structures
<code>void prefetch_access(AccessStat stat)</code>	Notifies the prefetcher about a cache access
<code>void prefetch_complete(Addr addr)</code>	Notifies the prefetcher about a prefetch load that has just completed

Table 3.2: Functions called by the simulator.

Function	Description
<code>void issue_prefetch(Addr addr)</code>	Called by the prefetcher to initiate a prefetch
<code>int get_prefetch_bit(Addr addr)</code>	Is the prefetch bit set for <code>addr</code> ?
<code>int set_prefetch_bit(Addr addr)</code>	Set the prefetch bit for <code>addr</code>
<code>int clear_prefetch_bit(Addr addr)</code>	Clear the prefetch bit for <code>addr</code>
<code>int in_cache(Addr addr)</code>	Is <code>addr</code> currently in the L2 cache?
<code>int in_mshr_queue(Addr addr)</code>	Is there a prefetch request for <code>addr</code> in the MSHR (miss status holding register) queue?
<code>int current_queue_size(void)</code>	Returns the number of queued prefetch requests
<code>void DPRINTF(trace, format, ...)</code>	Macro to print debug information. <code>trace</code> is a trace flag (<code>HWPrefetch</code>), and <code>format</code> is a <code>printf</code> format string.

Table 3.3: Functions callable from the user-defined prefetcher.

AccessStat member	Description
<code>Addr pc</code>	The address of the instruction that caused the access (Program Counter)
<code>Addr mem_addr</code>	The memory address that was requested
<code>Tick time</code>	The simulator time cycle when the request was sent
<code>int miss</code>	Whether this demand access was a cache hit or miss

Table 3.4: `AccessStat` members.

The prefetcher must implement the three functions `prefetch_init`, `prefetch_access` and `prefetch_complete`. The implementation may be empty.

The function `prefetch_init(void)` is called at the start of the simulation to allow the prefetcher to initialize any data structures it will need.

When the L2 cache is accessed by the CPU (through the L1 cache), the function `void prefetch_access(AccessStat stat)` is called with an argument (`AccessStat stat`) that gives various information about the access.

When the prefetcher decides to issue a prefetch request, it should call `issue_prefetch(Addr addr)`, which queues up a prefetch request for the block containing `addr`.

When a cache block that was requested by `issue_prefetch` arrives from memory, `prefetch_complete` is called with the address of the completed request as parameter.

Prefetches issued by `issue_prefetch(Addr addr)` go into a prefetch request queue. The cache will issue requests from the queue when it is not fetching data for the CPU. This queue has a fixed size (available as `MAX_QUEUE_SIZE`), and when it gets full, the oldest entry is evicted. If you want to check the current size of this queue, use the function `current_queue_size(void)`.

3.3 Using the interface

Start by studying `interface.hh`. This is the only M5-specific header file you need to include in your header file. You might want to include standard header files for things like printing debug information and memory allocation. Have a look at what the supplied example prefetcher (a very simple sequential prefetcher) to see what it does.

If your prefetcher needs to initialize something, `prefetch_init` is the place to do so. If not, just leave the implementation empty.

You will need to implement the `prefetch_access` function, which the cache calls when accessed by the CPU. This function takes an argument, `AccessStat stat`, which supplies information from the cache: the address of the executing instruction that accessed cache, what memory address was accessed, the cycle tick number, and whether the access was a cache miss. The block size is available as `BLOCK_SIZE`. Note that you probably will not need all of this information for a specific prefetching algorithm.

If your algorithm decides to issue a prefetch request, it must call the `issue_prefetch` function with the address to prefetch from as argument. The cache block containing this address is then added to the prefetch request

queue. This queue has a fixed limit of `MAX_QUEUE_SIZE` pending prefetch requests. Unless your prefetcher is using a high degree of prefetching, the number of outstanding prefetches will stay well below this limit.

Every time the cache has loaded a block requested by the prefetcher, `prefetch_complete` is called with the address of the loaded block.

Other functionality available through the interface are the functions for getting, setting and clearing the prefetch bit. Each cache block has one such tag bit. You are free to use this bit as you see fit in your algorithms. Note that this bit is *not* automatically set if the block has been prefetched, it has to be set manually by calling `set_prefetch_bit`. `set_prefetch_bit` on an address that is not in cache has no effect, and `get_prefetch_bit` on an address that is not in cache will always return false.

When you are ready to write code for your prefetching algorithm of choice, put it in `prefetcher/prefetcher.cc`. When you have several prefetchers, you may want to make `prefetcher.cc` a symlink.

The prefetcher is statically compiled into M5. After `prefetcher.cc` has been changed, recompile with `./compile.sh`. No options needed.

3.3.1 Example prefetcher

```
/*
 * A sample prefetcher which does sequential one-block lookahead.
 * This means that the prefetcher fetches the next block after_ the one that
 * was just accessed. It also ignores requests to blocks already in the cache.
 */

#include "interface.hh"

void prefetch_init(void)
{
    /* Called before any calls to prefetch_access. */
    /* This is the place to initialize data structures. */

    DPRINTF(HWPrefetch, "Initialized sequential-on-access prefetcher\n");
}

void prefetch_access(AccessStat stat)
{
    /* pf_addr is now an address within the next_ cache block */
    Addr pf_addr = stat.mem_addr + BLOCK_SIZE;

    /*
     * Issue a prefetch request if a demand miss occurred,
     * and the block is not already in cache.
     */
    if (stat.miss && !in_cache(pf_addr)) {
        issue_prefetch(pf_addr);
    }
}

void prefetch_complete(Addr addr) {
    /*
     * Called when a block requested by the prefetcher has been loaded.
     */
}
```

Chapter 4

Statistics

This chapter gives an overview of the statistics by which your prefetcher is measured and ranked.

IPC instructions per cycle. Since we are using a superscalar architecture, IPC rates > 1 is possible.

Speedup Speedup is a commonly used proxy for overall performance when running benchmark tests suites.

$$\text{speedup} = \frac{\text{execution time}_{\text{no prefetcher}}}{\text{execution time}_{\text{with prefetcher}}} = \frac{IPC_{\text{with prefetcher}}}{IPC_{\text{no prefetcher}}}$$

Good prefetch The prefetched block is referenced by the application before it is replaced.

Bad prefetch The prefetched block is replaced without being referenced.

Accuracy Accuracy measures the number of useful prefetches issued by the prefetcher.

$$acc = \frac{\text{good prefetches}}{\text{total prefetches}}$$

Coverage How many of the potential candidates for prefetches were actually identified by the prefetcher?

$$cov = \frac{\text{good prefetches}}{\text{cache misses without prefetching}}$$

Identified Number of prefetches generated and queued by the prefetcher.

Issued Number of prefetches issued by the cache controller. This can be significantly less than the number of *identified* prefetches, due to duplicate prefetches already found in the prefetch queue, duplicate prefetches found in the MSHR queue, and prefetches dropped due to a full prefetch queue.

Misses Total number of L2 cache misses.

Degree of prefetching Number of blocks fetched from memory in a single prefetch request.

Harmonic mean A kind of average used to aggregate each benchmark speedup score into a final average speedup.

$$H_{avg} = \frac{n}{\frac{1}{x_1} + \frac{1}{x_2} + \dots + \frac{1}{x_n}} = \frac{n}{\sum_{i=1}^n \frac{1}{x_i}}$$

Chapter 5

Debugging the prefetcher

5.1 m5.debug and trace flags

When debugging M5 it is best to use binaries built with debugging support (m5.debug), instead of the standard build (m5.opt). So let us start by recompiling M5 to be better suited to debugging:

```
scons -j2 ./build/ALPHA_SE/m5.debug.
```

To see in detail what's going on inside M5, one can specify enable trace flags, which selectively enables output from specific parts of M5. The most useful flag when debugging a prefetcher is `HWPrefetch`. Pass the option `--trace-flags=HWPrefetch` to M5:

```
./build/ALPHA_SE/m5.debug --trace-flags=HWPrefetch [...]
```

Warning: this can produce a lot of output! It might be better to redirect `stdout` to file when running with `--trace-flags` enabled.

5.2 GDB

The GNU Project Debugger `gdb` can be used to inspect the state of the simulator while running, and to investigate the cause of a crash. Pass GDB the executable you want to debug when starting it.

```
gdb --args m5/build/ALPHA_SE/m5.debug --remote-gdb-port=0  
-re --outdir=output/ammp-user m5/configs/example/se.py  
--checkpoint-dir=lib/cp --checkpoint-restore=1000000000  
--at-instruction --caches --l2cache --standard-switch  
--warmup-insts=10000000 --max-inst=10000000 --l2size=1MB  
--bench=ammp --prefetcher=on_access=true:policy=proxy
```

You can then use the `run` command to start the executable.

Some useful GDB commands:

run <args>	Restart the executable with the given command line arguments.
run	Restart the executable with the same arguments as last time.
where	Show stack trace.
up	Move up stack trace.
down	Move down stack frame.
print <expr>	Print the value of an expression.
help	Get help for commands.
quit	Exit GDB.

GDB has many other useful features, for more information you can consult the GDB User Manual at <http://sourceware.org/gdb/current/onlinedocs/gdb/>.

5.3 Valgrind

Valgrind is a very useful tool for memory debugging and memory leak detection. If your prefetcher causes M5 to crash or behave strangely, it is useful to run it under Valgrind and see if it reports any potential problems.

By default, M5 uses a custom memory allocator instead of `malloc`. This will not work with Valgrind, since it replaces `malloc` with its own custom memory allocator. Fortunately, M5 can be recompiled with `NO_FAST_ALLOC=True` to use normal `malloc`:

```
scons NO_FAST_ALLOC=True ./m5/build/ALPHA_SE/m5.debug
```

To avoid spurious warnings by Valgrind, it can be fed a file with warning suppressions. To run M5 under Valgrind, use

```
valgrind --suppressions=lib/valgrind.suppressions  
./m5/build/ALPHA_SE/m5.debug [...]
```

Note that everything runs *much* slower under Valgrind.